

Extraction of Salient Features for Mobile Robot Navigation via Teleoperation

Jian Peng
Dept. of Electrical Engineering
Tennessee State University
Nashville, TN 37209
jian.peng@ieee.org

Alan Peters
Dept. of Electrical Engineering and Computer Science
Vanderbilt University
Nashville, TN 37235
Rap2@vuse.vanderbilt.edu

Abstract

This paper presents a method to extract salient features from sensory-motor sequences for mobile robot navigation via teleoperation. Salient feature extraction consists of three steps: teleoperation, offline association, and evaluation. First, the mobile robot is teleoperated in an environment along a path several times. All sensory data and motor drive commands, are recorded. During an offline association step, these sensory-motor sequences are partitioned into episodes according to changes in motor commands. Salient features are then extracted by using two statistical criteria: consistency and correlation with the motor commands within the episode boundaries. Finally, these features are used to drive the robot in the learned environment. Some experiment results are also presented.

I. INTRODUCTION

MODERN mobile robots are usually equipped with a suite of different sensors [1]. To be successful in task execution the robot must detect features in the sensory data stream that are salient to the task. Behavior-based robots [2] in particular use salient features to construct and sequence behaviors. Salient feature extraction is, however, a notoriously difficult problem. In part this is due to the high dimensionality of the feature set in, for example, computer vision. [3]. In mid 1990's Pfeifer introduced the sensory-motor coordination (SMC) principle in [4] where he showed that the coupling of salient feature detection with motor control can reduce the dimensionality search space significantly. The relative simplicity of detecting changes in motor behavior, simplifies the search for salient features which, as Pfeifer demonstrated, tend to co-occur within a short period of time. The work of other researchers, notably Cohen [5] and Grupen [6], has supported Pfeifer's conclusions.

Those features that are salient to a specific task can only be determined *ex post facto*, after the robot has performed the task a number of times. Hence, the process must be bootstrapped. This paper reports on the use of teleoperation to do so. The idea is that if a person drives a robot through an environment making turns, stops, and starts, it is the person's sensory perceptions that trigger the changes in motor behavior. Through its own sensors, the robot might

pickup cues to the behavior changes effected by the teleoperator. If the task is repeated several times, some cues are likely to be present in all the trials. Those can be detected later through offline analysis and used to train the robot to perform a behavior sequence that completes the task in the same environment.

Such is the premise of this work. One must recall, however, that correlation does not imply causation, so some of the sensory cues that consistently co-occur with a behavior change may, in fact, be spurious. Nevertheless, we presumed that at a number of sensory cues would prove to be causal for each behavior change and that those cues would be salient – sufficient to effect the appropriate behavior change at the appropriate location and time for the robot to be successful in a navigation task. The results presented here support that contention. The paper is organized as follows. Section II introduces the sensory-motor coordination principle. Section III presents the salient feature extraction process. Section IV discusses the experimental results, and Section V concludes the paper.

II. SENSORY-MOTOR COORDINATION

Psychology has shown that perception and action are not two separate processes; instead, they are tightly coupled. This sensory-motor coordination [4] causes object-related actions to structure the sensory space. It is an active process whereby the agent's motions constrain its sensory input.

We define an SMC *event* as a motor action that occurs in response to a sensory stimulus, or a sensory stimulus that occurs in response to a motor action. We define an SMC *episode* as the state of the robot's sensors and actuators between consecutive SMC events. One can infer from Pfeifer's results that if SMC episodes are stored as vectors, clusters form in this vector space. The clusters, in effect, form a categorization of the robot's world – how its actions structure its sensing and vice-versa.

Our ideas for learning task-level behavior sequences have been based on a number of assumptions. When a teleoperator performs a task it is her/his SMC that is controlling the robot. So controlled, the robot's sensors detect its own internal states and those of the environment as it moves within it. Thus the robot can make its own associations between coincident motor actions and sensory

features as it is teleoperated. In repeating a task several times, a teleoperator will perform similar (albeit non-identical) sequences of motor actions whose dynamics will depend on her/his perceptions of similar sensory events that occur in similar sequence. As a result, the robot will detect a similar set of SMC events during each trial. Therefore, each trial can be partitioned into different episodes, demarcated by the common SMC events. We assume that sensory events that are salient to the task will occur in every trial and that sensory signals that differ across trials are insignificant to the task and can be ignored. By simply averaging the time-series for each episode point-wise over the trials, a canonical representation of the motor control sequence can be constructed. This motor sequence can be used to navigate the robot in the same environment later. This approach has been applied on NASA's Robonaut for the task learning [7]. It showed that six trials of a reach-grasp-release-retract skill are sufficient for learning a canonical description and construction of similar behaviors.

III. SALIENT FEATURE EXTRACTION

A. Procedure to Extract Salient Features

Our salient feature detection for mobile robot navigation was implemented in three steps: teleoperation and data recording, off-line data association, and execution and evaluation. First, the robot was teleoperated within an environment along a path. This was done several times. Each time, the robot followed similar paths with some deviations. All sensory data, both visual and non-visual, along with the motor commands were saved for the whole trial.

In the off-line association phase, signal analysis techniques constrained by motor commands were used (cf. next section) to identify salient features that were relevant to the task. These techniques also identified the events that consistently co-occurred with specific motor states or state transitions. Such a co-occurrence is assumed to be a SMC event – a motor state together with sensory pre-conditions and post-conditions. Combining these events forms a chain of behaviors.

Finally, the robot was put back into the same environment somewhere along the path over which it had been teleoperated previously. It was commanded to go to another position on the path. We found that the robot would consistently navigate to the target position using the learned features and SMC events

B. Off-line Data Association

1) Motor Commands and Episodes

The key step in the procedure is the off-line association phase, which involves several steps. First, the motor command sequences are used to partition the sensory-motor

time-series into *episodes*. A number of techniques for this have been proposed [8, 9]. Since the mobile robots used here have only three degrees-of-freedom, their motions are relatively simple. In our experiments, episodes were partitioned according to the motor command variations, i.e., a significant change in motor command demarcates an episode boundary. A simple thresholding of the motor velocities detects these boundaries.

After they are segmented, the episodes are time-normalized. Assume there were M trials, T_1, \dots, T_M , of the task performed via teleoperation. For each trial, T_i , assume N separate signals (from different sensors) $\mathbf{s}_{i,j}(t)$, were recorded. Then

$$\mathbf{v}_i(t) = \begin{bmatrix} \mathbf{s}_{i,1} & \dots & \mathbf{s}_{i,N} \end{bmatrix}^T(t) \quad (1)$$

is the vector time-series recorded during trial T_i . In general, $\mathbf{s}_{i,j}(t)$ is itself a vector time-series, such as odometry $[x, y, \theta]^T$. But $\mathbf{s}_{i,j}(t)$ could also be a scalar signal, such as digital compass output. We assume that t is discrete, defined only at integer multiples of the sampling interval, τ , so that

$$t \in \{n\tau\}_{n=1}^{\infty} \quad (2)$$

Hence, without loss of generality we define $t \in \mathbb{Z}^+$, the positive integers.

By assumption, each trial contains the same number, P , of episodes, $E_{i,k}$, which follow the same sequence, $E_{i,1}, \dots, E_{i,P}$, within each trial. Thus $\mathbf{s}_{i,j}(t)$ is the j -th signal from the i -th trial, and $E_{i,k}$ is the k -th episode from the i -th trial. Moreover,

$$E_{i,k} = \{\mathbf{v}_i(t)\}_{t=(t_{i,k-1})+1}^{t_{i,k}} \quad (3)$$

for $k = \{1, \dots, P\}$, where $t_{i,k-1} + 1$ is the time at which the k -th episode, $E_{i,k}$, starts in trial T_i , and $t_{i,k}$ is the time at which the k -th episode, $E_{i,k}$, ends. We define $t_{i,0} + 1$ as the starting time of trial T_i . Note that in general

$$t_{\eta,k} - t_{\eta,k-1} \neq t_{\nu,k} - t_{\nu,k-1} \quad (4)$$

That is the k -th episode from trial η will not have the same duration as the k -th episode from trial ν .

If the tasks were all recorded with the same τ , then the number of samples in corresponding episodes will differ. Let $\#\{\bullet\}$ represent the cardinality operator so that $\#\{E_{i,k}\}$ is the number of samples in episode k of trial i . Then usually

$$\#\{E_{\eta,k}\} \neq \#\{E_{\nu,k}\} \text{ if } \eta \neq \nu \quad (5)$$

To compute a characteristic representation of the task, the corresponding episodes in each task must have the same duration – the same number of samples. Therefore, each episode, $E_{i,k}$, was resampled to form a new one, $E'_{\eta,k}$ such

that

$$\#\{E'_{1,k}\} = \dots = \#\{E'_{M,k}\} = \frac{1}{M} \sum_{i=1}^M \#\{E_{i,k}\} \quad (6)$$

The length of episode $E'_{i,k}$ is the average over all trials of the number of samples in k -th episode. And,

$$E'_{i,k} = \{\mathbf{v}'_i(t)\}_{t=i_{k-1}+1}^{i_k} \quad (7)$$

where $\mathbf{v}'_i(t)$ is the resampled vector time-series,

$$\mathbf{v}'_i(t) = [\mathbf{s}'_{i,1} \ \dots \ \mathbf{s}'_{i,N}]^T(t) \quad (8)$$

and indices $\{t_{i,k}\}_{k=1}^P$ have been reassigned to the new time-series.

Within each episode, each time-normalized signal was averaged across trials:

$$\bar{\mathbf{v}}_k(t) = \frac{1}{M} \sum_{i=1}^M \begin{bmatrix} \mathbf{s}'_{i,1} \\ \vdots \\ \mathbf{s}'_{i,N} \end{bmatrix}(t) \quad (9)$$

for each $k \in \{1, \dots, P\}$. The resultant characteristic time series is the concatenation of the episode sample-wise means:

$$\bar{\mathbf{v}}(t) = \{\bar{\mathbf{v}}_1(t) | \bar{\mathbf{v}}_2(t) | \dots | \bar{\mathbf{v}}_P(t)\} \quad (10)$$

2) Salient Feature Selection

The characteristic vector time series, $\bar{\mathbf{v}}(t)$, contains not only motor control data but also sensory signals. A simple replay of the average of the motor control signals is not very useful since the task could only be successful if the environment were in exactly the same configuration as it was when the data were collected. Instead, the sensor data were analyzed for features that occurred within a short time window of episode boundaries. We assume that any sensory event that was a random occurrence in the M trials would be diminished, cancelled, or masked by noise through the episode averaging procedure. On the other hand, any event that occurred consistently with respect to an episode boundary would be preserved by the averaging. Those features were selected as salient for that specific episode.

There is no universal definition of ‘‘saliency’’. From the above discussion it can be concluded that, for a feature to be regarded as ‘‘salient’’ for a behavior, the feature must satisfy at least the following two conditions:

1. It should be consistent between different trials;
2. It should correlate well with the behavior change within a narrow time window.

These two criteria are applied in the following salient feature selection procedure:

1. Consistency Criterion

A square window, W_k , is used on every episode

boundary B_k , and a *consistency index*, $CI_{j,k}$, is calculated for each signal by summing and averaging in the window across all the trials $i = 1, \dots, M$:

$$CI_{j,k} = \sum_{W_k} \left(\frac{1}{M} \sum_i \left| \frac{\mathbf{s}'_{i,j}(t) - \bar{\mathbf{s}}'_{i,j}(t)}{\max(\bar{\mathbf{s}}'_{i,j}(t))} \right| \right) \quad (11)$$

where $\mathbf{s}'_{i,j}(t)$ is the episode-normalized signal of trial i and feature j ; $\bar{\mathbf{s}}'_{i,j}(t)$ is the average signal across all the trials obtained by the procedure given above. The smaller the value of $CI_{j,k}$, the more consistent the $\mathbf{s}'_{i,j}(t)$ (and presumably $\mathbf{s}_{i,j}(t)$, the original, not resampled sensory signal) are across all the trials. If $CI_{j,k}$ is less than a preset threshold, feature j of signal $\mathbf{s}_{i,j}(t)$ is selected as a candidate salient feature for episode k .

Alternatively, a variance index, $VI_{j,k}$, can be calculated for window W_k :

$$VI_{j,k} = \sum_{W_k} \left(\sqrt{\frac{1}{M} \sum_i \left(\frac{\mathbf{s}'_{i,j}(t) - \bar{\mathbf{s}}'_{i,j}(t)}{\max(\bar{\mathbf{s}}'_{i,j}(t))} \right)^2} \right) \quad (12)$$

Similarly, the smaller the value of $VI_{j,k}$, the more consistent the $\mathbf{s}'_{i,j}(t)$ (and $\mathbf{s}_{i,j}(t)$, the original sensory signal) are across all the trials. The difference between these two indexes is that $CI_{j,k}$ uses H^1 distance while $VI_{j,k}$ uses H^2 distance. The question arises: what is the difference between these two measurements in real application? An answer is suggested empirically from an analysis of the experimental data.

2. Correlation Criterion

We compute the correlation of every sensory feature with the motor drive command sequences within every episode boundary W_k :

$$\rho_{s_j, m, W_k}(z) = E_{W_k} \{ (s_i(t+z) - \mu_{s_i, W_k})(m(t) - \mu_{m, W_k}) \} \quad (13)$$

Function $\rho_{s_j, m, W_k}(z)$ is the correlation coefficient between the feature sequence $s_i(t)$ and the motor sequence $m(t)$ within k -th boundary window W_k with a lag of z steps; $E_{W_k} \{\bullet\}$ is the expectation in the window W_k ; \cdot . Values

μ_{s_i, W_k} and μ_{m, W_k} are the mean values of $s_i(t)$ and $m(t)$ in the W_k window, respectively. We assume that the higher the correlation, the more salient the feature is for the behavior. All the features with a correlation higher than a preset threshold and a consistency index (or variance index) less than a preset threshold are chosen as the salient features for that behavior. If a salient feature happens before the motor command change, we consider it to be a sensory

precursor to the behavior; otherwise, it is a sensory response.

3) Event-Behavior Chain

The selected salient features and the episodes can be represented graphically as an event-behavior chain (Figure 1). In this chain, each SMC event is a directed edge and each node is an episode that corresponds to a behavior. During autonomous operation, the robot performs the task by comparing the current sensory reading with the salient features and decides the next behavior. If the robot starts at the first state from the training sequence, then it performs the task by following the feature-behavior sequence (providing that all the salient features were detected during the training phase and are then detected at runtime. As described, the approach is purely deterministic and, therefore, brittle. By chaining the behaviors on a multiply connected graph and linking alternative behaviors probabilistically, we believe that the results can be made robust.




Figure 1: SMC sequence represented as a chain

IV. EXPERIMENT RESULTS

A. Experiment Platform

All experiments were carried out on an ATRV-Jr mobile robot made by Real World Interface, Inc. The robot is controlled onboard by a standard AMD Athlon XP 1.4 GHz PC running Linux. . The sensors on the robot, include an odometer, 17 sonar sensors, a SICK laser scanner, an electrical compass, a Crossbow gyroscope called “DMU” (Dynamics Measurement Unit), and a Sony pan-tilt-zoom, monocular, color, active camera system. All low-level sensing and actuation modules run on the onboard computer. High-level modules and user interface are executed on a remote laptop computer.

B. Experimental Results

1) Teleoperation

Experiments were performed outdoors on sidewalks on campus at Vanderbilt University. The robot was teleoperated five times through a navigation task. The starting position and ending position were fixed. During each trial, all sensory information and motor commands were recorded at the rate of about 3 Hz. Each sequence had about 350 samples. Figure 2 shows images taken by the onboard camera at the starting and ending point. The trajectories of the trials are shown in Figure 3. The turns occurred at points where two or three sidewalks joined.




Figure 2: Images taken by the onboard camera during teleoperation

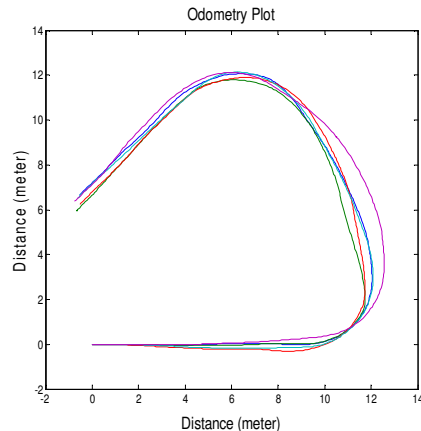


Figure 3: Robot trajectories of the five teleoperation trials according to the odometer.

2) Off-line Data Association

These trial sequences were partitioned into episodes and then time-normalized according to the method described in the previous section. Figure 4 shows the result, where dashed lines correspond to the five trials, while the solid line is their average. The plot suggests that there are four episode boundaries that demarcate five episodes: move forward, turn left, move-forward-while-turning-left, turn left, and move forward then stop. All sensory data were time-normalized according to these episodes. Figure 5 shows the time-normalized length of the longest segment in laser readings.

The consistency index and correlation index criteria were then used to select salient features at these episode boundaries. shows the result for the first episode boundary. By thresholding, the following features were selected as salient for each episode boundaries:

- Boundary 1: length of the longest continuous segment from the laser scanner (will be called LLSeg hereafter), DMU^1 heading, hue (0-0.1) count².

¹ DMU stands for “dynamics measurement unit”, a gyroscope that measures linear and angular acceleration along the three primal axes.

² Each image captured by the onboard camera is converted from RGB to HSV (Hue, Saturation, and Value) space. The ranges for H, S, and V are between 0 and 1. Feature *Hue (0-0.1) count* is the number of pixels in the image that have a hue value between 0 and 0.1. Other hue count features are similar.

- Boundary 2: LLSeg, compass, DMU heading, hue (0.2-0.3) count.
- Boundary 3: compass, DMU heading, hue (0-0.1) count, hue (0.9-1.0) count
- Boundary 4: LLSeg, compass, DMU heading.

An event-behavior chain was then constructed using these selected salient features.

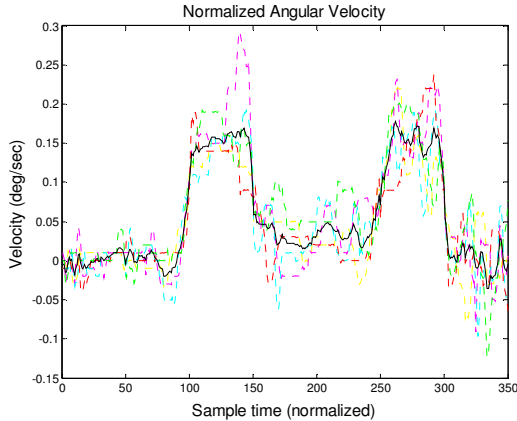


Figure 4: Time-normalized angular velocity profiles. The dashed lines are of the trials, while the solid line is the average.

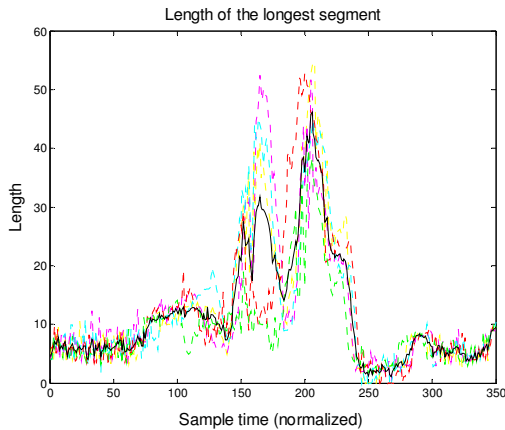


Figure 5: Time-normalized length of the longest segment from the laser sensor.

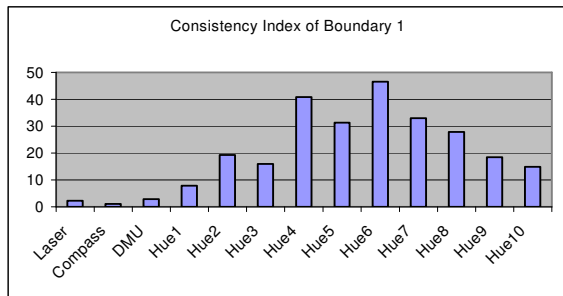


Figure 6: Consistency index of the boundary 1 event

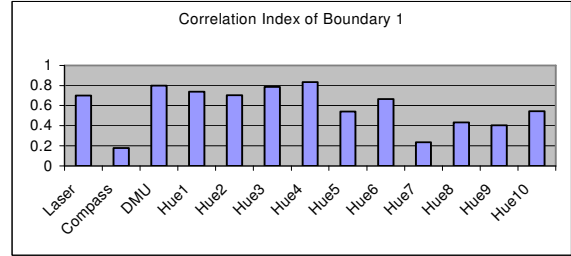


Figure 7: Correlation index of the boundary 1 event

3) Evaluation

As mentioned before, there is no universal definition for “saliency”, thus there is no universally accepted method to evaluate saliency. To evaluate the effectiveness of the selected salient features, the robot was put back in the environment where it was previously teleoperated and was commanded to go to the original target position based on the constructed event-behavior chain. The success rate is an indication of the effectiveness of the selected salient features.

In total, 15 evaluation navigation tasks were performed under different day, time, and weather conditions. 12 times the robot reached a destination within 1.5 meter of the target position³, thus the success rate was 80%. Figure 8 shows the trajectory plot of these successful trials, and Figure 9 shows the trajectory plot of the failed trials. Note that all the failures happened at Boundary 3.

C. Two Questions

Two questions arise naturally from the above discussion and experiments:

1. Two consistency indexes are proposed in Section III (B). Is there any significant difference between them in this experiment?
2. Five teleoperation trials were used to extract the salient features, and obtained satisfactory result. The question is: why five? How about 2, 3, or 4 trials?

These issues will be addressed empirically using the experimental data. For the first question, the two consistency indexes (Eq. 11 and 12) are based on H^1 and H^2 distance, and they produce compatible results. This is not surprising, since H^1 and H^2 are similar distance measurements.

To check the effect of the number of teleoperation trials on salient feature extraction, the original five teleoperation trials were combined into groups of two trials, three trials, and four trials. For group of two, three, and four trials, there were totally ten ($C_5^2 = 10$), ten ($C_5^3 = 10$), and five ($C_5^4 = 5$) combinations, respectively. Only consistency

³ The robot’s main body is about 0.4x0.8 meters. If tires and bumpers included, the size is about 0.7x1.0 meters. Comparing to this size, 1.5 meters is a reasonable measurement to determine success/failure.

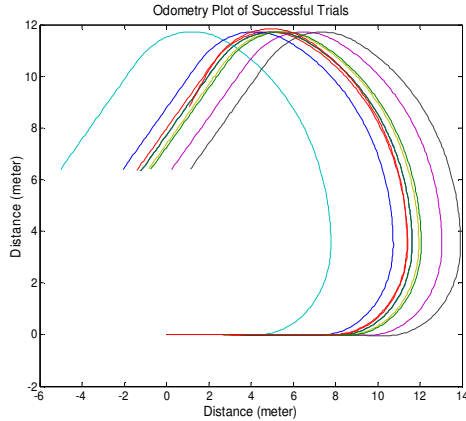


Figure 8: Trajectory plot of the 12 successful evaluation trials according to the odometer. Note that the odometer was reset at the beginning of every trial, and the robot started at different locations for each trial.

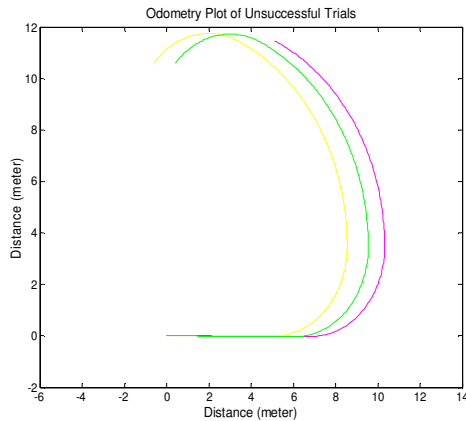


Figure 9: Trajectory plot of the 3 failed evaluation trials according to the odometer. Note that all the failures happened at Boundary 3.

Table 1: Error rate for different numbers of trials

	Type-I error rate	Type-II error rate
2 trials	21%	11%
3 trials	18%	9.2%
4 trials	5.4%	4.3%

criterion was considered here, because the other criterion, correlation, would remain similar irrelevant to the number of the trials.

Both type-I error (regard a random signal as a salient features) and type-II error (regard a salient feature as a random signal) were obvious in these combinations. Using the same threshold as that used for five teleoperation trials, the error rate is shown the following table: In our experiment, 5 trials gave satisfactory result. Another question arises: How about more trial (e.g., 6, 7, 8)? How many trials are necessary to extract salient features? These questions need to be explored in the future research.

V. CONCLUSION

A method to extract salient features in a specific environment for a specific navigation task is presented in this paper. Salient feature in this context means that a) it is consistent between different trials, and b) it correlates well with the behavior change at the episode boundary. This is a complement to what is known as novelty detection [10], where the goal is to recognize features that differ from those that are normally seen.

Based on these two criteria, a three-step method to extract salient features was presented. Speaking loosely, it is a supervised learning method. In [10], a self-organizing map (SOM) based on habituation was used to learn the environment and to discover novel features. In future research, a similar approach could be applied and could possibly detect salient feature extraction on-line, and thus no off-line association is necessary.

Initial experiments performed both indoors (not discussed here) and outdoors have shown that this is a promising approach. We are also extending this approach into feature detection for humanoid robots and other applications.

REFERENCES

- [1] H. R. Everett, *Sensors for Mobile Robots: Theory and Application*. Natick, MA: A K Peters, 1995.
- [2] R. C. Arkin, *Behavior-Based Robotics*. Cambridge, MA, USA: MIT Press, 1998.
- [3] G. N. DeSouza and A. C. Kak, "Vision for Mobile Robot Navigation: A Survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, pp. 237-267, 2002.
- [4] R. Pfeifer and C. Scheier, "Sensory-motor coordination: The metaphor and beyond," *Robotics and Autonomous Systems*, vol. 20, pp. 157-178, 1997.
- [5] P. R. Cohen, N. Adams, and D. Hand, "Finding Patterns that Correspond to Episodes," University of Massachusetts, Amherst, MA, Technical Report 01-11, 2001.
- [6] J. A. Coelho, Jr., J. H. Piater, and R. A. Grupen, "Developing Haptic and Visual Perceptual Categories for Reaching and Grasping with a Humanoid Robot," *Robotics and Autonomous Systems*, vol. 37, pp. 195-219, 2001.
- [7] R. A. Peters, C. Campbell, W. Bluethmann, and E. Huber, "Robonaut Task Learning through Teleoperation," presented at International Conference on Robotics and Automation, 2003.
- [8] P. R. Cohen and N. Adams, "An algorithm for segmenting categorical time series into meaningful episodes," presented at Fourth Symposium on Intelligent Data Analysis, 2001.
- [9] A. Fod, M. J. Mataric, and O. C. Jenkins, "Automated derivation of primitives for movement classification," *Autonomous Robots*, vol. 12, pp. 39-54, 2002.
- [10] S. Marsland, U. Nehmzow, and J. Shapiro, "On-line Novelty Detection for Autonomous Mobile Robots," *Journal of Robotics and Autonomous Systems*, 2004.